Lab 1

Interfacing With Other Programming Languages Using Cython

Lab Objective: *Learn to interface with object files using Cython. This lab should be worked through on a machine that has already been configured to build Cython extensions using gcc or MinGW.*

Suppose you are writing a program in Python, but would like to call code written in another language. Perhaps this code has already been debugged and heavily optimized, so you do not want to simply re-implement the algorithm in Python or Cython. In technical terms, you want Python to *interface* (or communicate) with the other language. For example, NumPy's linear algebra functions call functions from LAPACK and BLAS, which are written in Fortran.

One way to have Python interface with C is to write a Cython *wrapper* for the C function. The wrapper is a Cython function that calls the C function. This is relatively easy to do because Cython compiles to C. From Python, you can call the wrapper, which calls the C function (see Figure TODO).

Wrapping with Cython: an overview

When you use Cython to wrap a C function, you just write a Cython function that calls the C function. To actually do this, we need to understand a little bit more about how C works.

After you write a program in C, you pass it to a compiler. The compiler turns your C code into machine code, i.e., instructions that your computer can read. The output of the compiler is an *object file*.

In our scenario, we have C code defining a single function that we would like to wrap in Cython. This C code has already been compiled to an object file. The protocol for calling the C code from Cython is similar to calling it from C:

1. we need to *include a header file* for the C code in our Cython code, and

2. we must *link to the object file* compiled from the C code when we compile the Cython code.

A *header* for a file contains the declaration of the function defined in the file. We include the header in the Cython code so that the compiler can check that the function is called with the right number and types of arguments. Then, when we tell the compiler to *link to* the object file, we simply tell it where to find the instructions defining the function declared in the header.

The Cython code will compile to a second object file. This object file defines a module that can be imported into Python.

Wrapping with Cython: an example

As an example, we will wrap the C function `ctridiag()` below. This function computes the solution to a tridiagonal system $A\mathbf{v} = \mathbf{x}$. Its parameters are pointers to four 1-D arrays, which satisfy the following:

- Arrays `a` and `c` have length `n-1` and contain the first subdiagonal and super-diagonal of A, respectively.

- Arrays `b` and `x` have length `n` and represent the main diagonal of A and $\mathbf{x}$, respectively.

The array `c` is used to store temporary values and `x` is transformed into the solution of the system.

```c
void ctridiag(double *a, double *b, double *c, double *x, int n){
    // Solve a tridiagonal system inplace.
    // Initialize temporary variable and index.
    int i;
    double temp;
    // Perform necessary computation in place.
    // Initial steps
    c[0] = c[0] / b[0];
    x[0] = x[0] / b[0];
    // Iterate down arrays.
    for (i=0; i<n-2; i++){
        temp = 1.0 / (b[i+1] - a[i] * c[i]);
        c[i+1] = c[i+1] * temp;
        x[i+1] = (x[i+1] - a[i] * x[i]) * temp;
        }
    // Perform last step.
    x[n-1] = (x[n-1] - a[n-2] * x[n-2]) / (b[n-1] - a[n-2] * c[n-2]);
    // Perform back substitution to finish constructing the solution.
    for (i=n-2; i>-1; i--){
        x[i] = x[i] - c[i] * x[i+1];
        }
    }
```

ctridiag.c

This terminal command tells the C-compiler gcc to compiles `ctridiag.c` to an object file `ctridiag.o`.

```
>>> gcc -fPIC -c ctridiag.c -o ctridiag.o
```

The `-fPIC` option is required when building code for a shared object. The `-c` flag prevents the compiler from raising an error even though `ctridiag.c` does not have a `main` function.

Write a header for `ctridiag.c`

The header file essentially contains a function declaration for `ctridiag()`. It tells Cython how to use the object file `ctridiag.o`.

```
1  extern void ctridiag(double* a, double* b, double* c, double* x, int n);
2  // This declaration specifies what arguments should be passed
   // to the function called `ctridiag.'
```

ctridiag.h

Write a Cython wrapper

Next we write a Cython file containing a function that "wraps" `ctridiag()`. This file must include the header we wrote in the previous step.

```
1  # Include the ctridiag() function as declared in ctridiag.h
2  # We use a cdef since this function will only be called from Cython
   cdef extern from "ctridiag.h":
4      void ctridiag(double* a, double* b, double* c, double* x, int n)

6  # Define a Cython wrapper for ctridiag().
   # Accept four NumPy arrays of doubles
8  # This will not check for out of bounds memory accesses.
   cpdef cytridiag(double[:] a, double[:] b, double[:] c, double[:] x):
10     cdef int n = x.size
       ctridiag(&a[0], &b[0], &c[0], &x[0], n)
```

cython_ctridiag.pyx

Some comments on this code are in order. First, including the header for `ctridiag()` allows us to call this function even though it was not defined in this file. This is a little like importing NumPy at the start of your Python script.

Second, the arguments of the Cython function `cytridiag()` are not in bijection with the arguments of `ctridiag()`. In fact, Cython does not need to know the size of the NumPy arrays it accepts as inputs because these arrays are objects that carry that information with them. We extract this information and pass it to the C function inside the Cython wrapper.

However, it is possible to unintentionally access memory outside of an array by calling this function. For example, suppose the input arrays `a`, `b`, `c`, and `x` are of sizes 4, 5, 3, and 5. Since `x` is used to determine the size of the system, the function `ctridiag()` will expect `c` to have length 4. At some point, `ctridiag()` will likely try to read or even write to the 4th entry of `c`. This address does exist in memory, but it does not contain an entry of `c`! Therefore, this function must be called by a responsible user who knows the sizes of her arrays. Alternatively, you could check the sizes of the arrays in the Cython wrapper before passing them to the C function.

Finally, the C function expects a parameter `double* a`, meaning that `a` is a *pointer to* (i.e., the address of) a `double`. The funtion `ctridiag()` expects this double to be the first entry of the array `a`. So instead of passing the object `a` to `ctridiag()`, we find the first entry of `a` with `a[0]`, and then take its address with the `&` operator.

Compile the Cython wrapper

Now we can compile `cython_ctridiag.pyx` to build the Python extension. The following setup file uses distutils to compile the Cython file, and may be run on Windows, Linux, or Macintosh-based computers. Notice that in line 28, we link to the existing object file `ctridiag.o`.

```python
# Import needed setup functions.
from distutils.core import setup
from distutils.extension import Extension
# This is part of Cython's interface for distutils.
from Cython.Distutils import build_ext
# We still need to include the directory
# containing the NumPy headers.
from numpy import get_include
# We still need to run one command via command line.
from os import system

# Compile the .o file we will be accessing.
# This is independent building the Python extension module.
shared_obj = "gcc ctridiag.c -fPIC -c -o ctridiag.o"
print shared_obj
system(shared_obj)

# Tell Python how to compile the extension.
ext_modules = [Extension(
                        # Module name:
                        "cython_ctridiag",
                        # Cython source file:
                        ["cython_ctridiag.pyx"],
                        # Other compile arguments
                        # This flag doesn't do much this time,
                        # but this is where it would go.
                        extra_compile_args=["-fPIC", "-O3"],
                        # Extra files to link to:
                        extra_link_args=["ctridiag.o"])]

# Build the extension.
setup(name = 'cython_ctridiag',
      cmdclass = {'build_ext': build_ext},
      # Include the directory with the NumPy headers when compiling.
      include_dirs = [get_include()],
      ext_modules = ext_modules)
```

ctridiag_setup_distutils.py

This setup file can be called from the command line with the following command.

```
python ctridiag_setup_distutils.py build_ext --inplace
```

The `--inplace` flag tells the script to compile the extension in the current directory. The appendix at the end of this lab contains setup files that build the Python extension by hand on various operating systems.

Test the Python extension

After running the setup file, you should have a Python module called `cython_cytridiag` that defines a function `cytridiag()`. You can import this module into IPython as

usual. However, if you modify `ctridiag()` and then try to recompile the extension, you may get an error if the module is currently in use. Hence, if you are frequently recompiling your extension, it is wise to test it with a script.

The following script tests the module `cython_cytridiag`.

```python
import numpy as np
from numpy.random import rand
# Import the new module.
from cython_ctridiag import cytridiag as ct

# Construct arguments to pass to the function.
n=10
a, b, c, x = rand(n-1), rand(n), rand(n-1), rand(n)
# Construct a matrix A to test that the result is correct.
A = np.zeros((n,n))
A.ravel()[A.shape[1]::A.shape[1]+1] = a
A.ravel()[::A.shape[1]+1] = b
A.ravel()[1::A.shape[1]+1] = c
# Store x so we can verify the algorithm returned the correct values
x_copy = x.copy()
# Call the function.
ct(a, b, c, x)
# Test to see if the result is correct.
# Print the result in the command line.
if np.absolute(A.dot(x) - x_copy).max() < 1E-12:
    print "Test Passed"
else:
    print "Test Failed"
```

<center>ctridiag_test.py</center>

Another example

The C function `cssor()` below implements the Successive Over Relaxation algorithm for solving Laplace's equation, which in two dimensions is

$$\frac{\partial^2 F}{\partial x^2} + \frac{\partial^2 F}{\partial y^2} = 0.$$

In the Successive Over Relaxation algorithm, an input array is iteratively modified until it converges to the solution $F(x, y)$.

```c
void cssor(double* U, int m, int n, double omega, double tol, int maxiters, int* ↩
    info){
    // info is passed as a pointer so the function can modify it as needed.
    // Temporary variables:
    // 'maxerr' is a temporary value.
    // It is used to determine when to stop iteration.
    // 'i', 'j', and 'k' are indices for the loops.
    // lcf and rcf will be precomputed values
    // used on the inside of the loops.
    double maxerr, temp, lcf, rcf;
    int i, j, k;
    lcf = 1.0 - omega;
    rcf = 0.25 * omega;
    for (k=0; k<maxiters; k++){
```

```
14          maxerr = 0.0;
            for (j=1; j<n-1; j++){
16              for (i=1; i<m-1; i++){
                    temp = U[i*n+j];
18                  U[i*n+j] = lcf * U[i*n+j] + rcf * (U[i*n+j-1] + U[i*n+j+1] + U[(i↩
                        -1)*n+j] + U[(i+1)*n+j]);
                    maxerr = fmax(fabs(U[i*n+j] - temp), maxerr);}}
20          // Break the outer loop if within
            // the desired tolerance.
22          if (maxerr < tol){break;}}
        // Here we have it set status to 0 if
24      // the desired tolerance was attained
        // within the the given maximum
26      // number of iterations.
        if (maxerr < tol){*info=0;}
28      else{*info=1;}}
```

cssor.c

The function `cssor()` accepts the following inputs.

- U is a pointer to a Fortran-contiguous 2-D array of double precision numbers.

- m is an integer storing the number of rows of U.

- n is an integer storing the number of columns of U.

- ω is a double precision floating point value between 1 and 2. The closer this value is to 2, the faster the algorithm will converge, but if it is too close the algorithm may not converge at all. For this lab just use 1.9.

- tol is a floating point number storing a tolerance used to determine when the algorithm should terminate.

- maxiters is an integer storing the maximum allowable number of iterations.

- info is a pointer to an integer. This function will set info to 0 if the algorithm converged to a solution within the given tolerance. It will set info to 1 if it did not.

Problem 1. Wrap `cssor()` so it can be called from Python, following these steps.

1. Write a C header for `cssor.c`.

2. Write a Cython wrapper `cyssor()` for the function `cssor()` as follows.

   (a) Have `cyssor()` accept parameters `tol` and `maxiters` that default to $10e^{-8}$ and 10,000, respectively. What other arguments does `cyssor ()` need to accept?

   (b) Check that the input `U` is a Fortran-contiguous array (this means that the array is stored in a certain way in memory). You can check if `U` is Fortran-contiguous with the code `U.flags["F_CONTIGUOUS`

`"]`, which returns a Boolean value. If `U` is not Fortran-contiguous, raise a ValueError as follows:

```
raise ValueError('Input array U is not Fortran-contiguous')
```

When raising a ValueError, you may replace the string `'Input array U is not Fortran-contiguous'` with any desired text. For more about errors and exceptions in Python, see `https://docs.python.org/2/tutorial/errors.html`.

(c) Raise a ValueError if `cssor()` fails to converge—i.e., if `cssor()` sets `info` to 1.

3. Compile `cssor.c` to an object file and compile your Cython wrapper to a Python extension. Be aware that you may compile `cssor.c` correctly and still receive warnings from the compiler.

4. Write a test script for your Cython function.

(a) You can run the function with the following code.

```
import numpy as np
from cython_fssor import cyssor
resolution = 501
U = np.zeros((resolution, resolution))
X = np.linspace(0, 1, resolution)
U[0] = np.sin(2 * np.pi * X)
U[-1] = - U[0]
cyssor(U, 1.9)
```

(b) Have your test script run the code above and plot the modified array `U` on the domain $[0, 1] \times [0, 1]$. Note that U is a 501×501 array. Your plot should look like Figure 1.1.
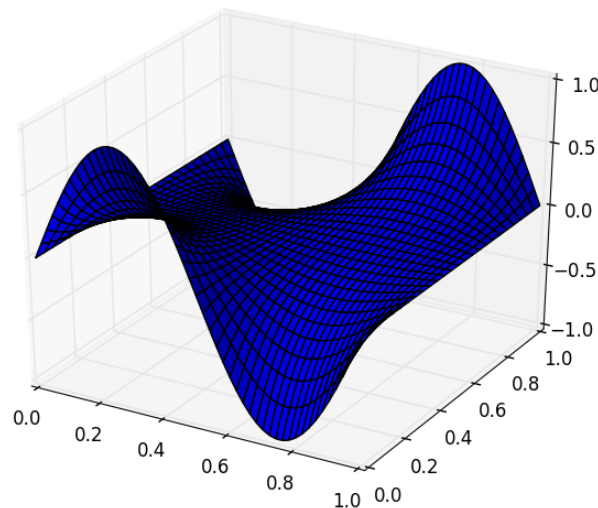
Figure 1.1: Correct output for the test script of Problem 1.

Wrapping a Fortran function (Optional)

We can also use Cython to wrap Fortran. As an example, we will wrap the Fortran function below, which implements the same algorithm as `ctridiag()`.

We have used the C library `iso_c_binding` to make the function accept pointers to native C types. If we were wrapping a function or subroutine that we did not write ourselves, we would have to define a Fortran function that uses the `iso_c_binding` library to accept pointers from C and then uses the values it receives to call the original function.

```fortran
module tridiag

use iso_c_binding, only: c_double, c_int

implicit none

contains

! The expression bind(c) tells the compiler to
! make the naming convention in the object file
! match the naming convention here.
! This will be a subroutine since it does not
! return any values.
subroutine ftridiag(a, b, c, x, n) bind(c)

!    Here we declare the types for the inputs.
!    Thisis where we use the c_double and c_int types.
```

```fortran
18  !   The 'dimension' statement tells the compiler that
    !   the argument is an array of the given shape.
20      integer(c_int), intent(in) :: n
        real(c_double),dimension(n),intent(in) :: b
22      real(c_double),dimension(n),intent(inout) :: x
        real(c_double),dimension(n-1),intent(in) :: a
24      real(c_double),dimension(n-1),intent(inout) :: c

26  !   Two temporary varaibles.
    !   'm' is a temporary value.
28  !   'i' is the index for the loops.
        real(c_double) m
30      integer i

32  !   Here is the actual computation:
        c(1) = c(1) / b(1)
34      x(1) = x(1) / b(1)
    !   This is the syntax for a 'for' loop in Fortran.
36  !   Indexing for arrays in fortran starts at 1
    !   instead of starting at 0 like it does in Python.
38  !   Arrays are accessed using parentheses
    !   instead of brackets.
40      do i = 1,n-2
            m = 1.0D0 / (b(i+1) - a(i) * c(i))
42          c(i+1) = c(i+1) * m
            x(i+1) = x(i+1) - a(i) * x(i)
44          x(i+1) = x(i+1) * m
    !   Note that you have to explicitly end the loop.
46      enddo
        x(n) = (x(n) - a(n-1) * x(n-1)) / (b(n) - a(n-1) * c(n-1))
48      do i = n-1,1,-1
            x(i) = x(i) - c(i) * x(i+1)
50      enddo
    !   You must also explicitly end the function or subroutine.
52  end subroutine ftridiag

54  end module
```

ftridiag.f90

> **WARNING**
>
> When interfacing between Fortran and C, you will have to pass pointers to *all* the variables you send to the Fortran function as arguments. Passing a variable directly will probably crash Python.

Write a header for `ftridiag.f90`

Here is a header that tells C how to interface with the function we have just defined.

```c
1  extern void ftridiag(double* a, double* b, double* c, double* x, int* n);
2  // Notice that we passed n as a pointer as well.
```

ftridiag.h

To compile `ftridiag.f90` to an object file you can run the following command in your command line:

```
gfortran -fPIC -c ftridiag.f90 -o ftridiag.o
```

Write a Cython wrapper, compile, and test it

The Cython wrapper for this function is analogous to `cython_ctridiag.pyx`, except that every variable passed to `ftridiag` should be a pointer. Aside from the use of gfortran instead of gcc, the rest of the compilation and testing process is entirely the same. The setup files specific to Windows and Linux, the setup file using distutils, and the test script are included in this lab folder.

> **WARNING**
>
> Fortran differs from C in that the columns of arrays are stored in contiguous blocks of memory instead of the rows. The default for NumPy and for C arrays is to have rows stored in contiguous blocks, but in NumPy this varies. You should make sure any code passing arrays between Fortran, C, and NumPy addresses this discrepancy. Though by default Fortran assumes arrays are arranged in Fortran-order, for many routines you can specify whether to act on an array or its transpose.

> **NOTE**
>
> The function `ftridiag()` may seem unusual because it is a Fortran function that is callable from C. In fact, this is not an uncommon scenario. Many larger Fortran libraries (for example, LAPACK) provide C wrappers for all their functions.

Another example

Here is a Fortran implementation of the Successive Over Relaxation algorithm for solving Laplace's equation. This function is the Fortran equivalent of `cssor.c`. Its parameters are the same, with the exception that all parameters are *points* to the values they reference.

```fortran
module ssor

use iso_c_binding, only: c_double, c_int

implicit none

contains

! The expression bind(c) tells the compiler to
```

```fortran
10  ! make the naming convention in the object file
    ! match the naming convention here.
12  ! This will be a subroutine since it does not
    ! return any values.
14  subroutine fssor(U, m, n, omega, tol, maxiters, info) bind(c)

16  !   Here we declare the types for the inputs.
    !   This is where we use the c_double and c_int types.
18  !   The 'dimension' statement tells the compiler that
    !   the argument is an array of the given shape.
20      integer(c_int), intent(in) :: m, n, maxiters
        integer(c_int), intent(out) :: info
22      real(c_double), dimension(m,n), intent(inout) :: U
        real(c_double), intent(in) :: omega, tol

24
    !   Temporary variables:
26  !   'maxerr' is a temporary value.
    !   It is used to determine when to stop iteration.
28  !   'i', 'j', and 'k' are indices for the loops.
    !   lcf and rcf will be precomputed values
30  !   used on the inside of the loops.
        real(c_double) :: maxerr, temp, lcf, rcf
32      integer i, j, k

34      lcf = 1.0D0 - omega
        rcf = 0.25D0 * omega
36      do k = 1, maxiters
            maxerr = 0.0D0
38          do j = 2, n-1
                do i = 2, m-1
40                  temp = U(i,j)
                    U(i,j) = lcf * U(i,j) + rcf * (U(i,j-1) + U(i,j+1) + U(i-1,j) + U↩
                      (i+1,j))
42                  maxerr = max(abs(U(i,j) - temp), maxerr)
                enddo
44          enddo
    !       Break the outer loop if within
46  !       the desired tolerance.
            if (maxerr < tol) exit
48      enddo
    !   Here we have it set status to 0 if
50  !   the desired tolerance was attained
    !   within the the given maximum
52  !   number of iterations.
        if (maxerr < tol) then
54          info = 0
        else
56          info = 1
        end if
58  end subroutine fssor

60  end module
```

fssor.f90

Problem 2. Imitate Problem 1 to wrap `fssor()` with Cython.

> **WARNING**
>
> When passing arrays as pointers to C, C++, and Fortran functions, be *absolutely sure* that the array being passed is contiguous. This means that the entries of the array are stored in *adjacent* entries in memory. Passing one of these functions a strided array will result in out of bounds memory accesses and could crash your computer. When working with two dimensional arrays, C and C++ expect rows to be stored in contiguous blocks of memory. Fortran expects columns to be stored in contiguous blocks of memory.
>
> To avoid out of bounds memory accesses, be sure to pass in the size of the array as a separate parameter.

Appendix: compiling C extensions for Python

If you know more about compilers, you may find it enlightening to manually compile a C extension for Python. Here we show how to manually compile `cython_ctridiag.pyx` on Windows, Linux, and Macintosh machines.

On Windows, using the compiler MinGW (a version of gcc for windows), the compilation can be performed by running the following setup file.

```python
# The system function is used to run commands
# as if they were run in the terminal.
from os import system
# Get the directory for the NumPy header files.
from numpy import get_include
# Get the directory for the Python header files.
from distutils.sysconfig import get_python_inc
# Get the directory for the general Python installation.
from sys import prefix

# Now we construct a string for the flags we need to pass to
# the compiler in order to give it access to the proper
# header files and to allow it to link to the proper
# object files.
# Use the -I flag to include the directory containing
# the Python headers for C extensions.
# This header is needed when compiling Cython-made C files.
# This flag will look something like: -Ic:/Python27/include
ic = " -I" + get_python_inc()
# Use the -I flag to include the directory for the NumPy
# headers that allow C to interface with NumPy arrays.
# This is necessary since we want the Cython file to operate
# on NumPy arrays.
npy = " -I" + get_include()
# This links to a compiled Python library.
# This flag is only necessary on Windows.
lb = " \"" + prefix + "/libs/libpython27.a\""
# This links to the actual object file itself
o = " \"ctridiag.o\""

# Make a string of all these flags together.
# Add another flag that tells the compiler to
# build for 64 bit Windows.
```

```python
34  flags = "-fPIC" + ic + npy + lb + o + " -D MS_WIN64"

36  # Build the object file from the C source code.
    system("gcc ctridiag.c -fPIC -c -o ctridiag.o")
38  # Compile the Cython file to C code.
    system("cython -a cython_ctridiag.pyx --embed")
40  # Compile the C code generated by Cython to an object file.
    system("gcc -c cython_ctridiag.c" + flags)
42  # Make a Python extension containing the compiled
    # object file from the Cython C code.
44  system("gcc -shared cython_ctridiag.o -o cython_ctridiag.pyd" + flags)
```

ctridiag_setup_windows64.py

The following file works on Linux and Macintosh machines using gcc.

```python
1   # The system function is used to run commands
2   # as if they were run in the terminal.
    from os import system
4   # Get the directory for the NumPy header files.
    from numpy import get_include
6   # Get the directory for the Python header files.
    from distutils.sysconfig import get_python_inc

8
    # Now we construct a string for the flags we need to pass to
10  # the compiler in order to give it access to the proper
    # header files and to allow it to link to the proper
12  # object files.
    # Use the -I flag to include the directory containing
14  # the Python headers for C extensions.
    # This header is needed when compiling Cython-made C files.
16  ic = " -I" + get_python_inc()
    # Use the -I flag to include the directory for the NumPy
18  # headers that allow C to interface with NumPy arrays.
    # This is necessary since we want the Cython file to operate
20  # on NumPy arrays.
    npy = " -I" + get_include()
22  # This links to the actual object file itself
    o = " ctridiag.o"

24
    # Make a string of all these flags together.
26  # Add another flag that tells the compiler to make
    # position independent code.
28  flags = ic + npy + o + " -fPIC"

30  # Build the object file from the C source code.
    system("gcc ctridiag.c -c -o ctridiag.o -fPIC")
32  # Compile the Cython file to C code.
    system("cython -a cython_ctridiag.pyx --embed")
34  # Compile the C code generated by Cython to an object file.
    system("gcc -c cython_ctridiag.c" + flags)
36  # Make a Python extension containing the compiled
    # object file from the Cython C code.
38  system("gcc -shared cython_ctridiag.o -o cython_ctridiag.so" + flags)
```

ctridiag_setup_linux.py